

REMARKS

The Examiner has rejected Claims 1, 2, 4, 11, 12, 14 and 21 under 35 U.S.C. 102(e) as being anticipated by Illikkal et al. (U.S. Patent Publication No. 2005/0066028). Applicant respectfully disagrees with such rejection.

With respect to independent Claims 1 and 11, the Examiner has relied on Fig. 2 from the Illikkal reference to make a prior art showing of applicant's claimed technique "wherein the CB cache is a multi-port device providing direct access to the CB cache via each port."

Applicant respectfully asserts that Fig. 2 from the Illikkal reference relied upon by the Examiner merely teaches that "[t]o reduce the burden on the host processor 210, a TCP/IP offload engine 270 may be provided on the NIC 250" (Paragraph [0017] – emphasis added) and that "[t]o avoid frequent and time-consuming accesses to the external memory unit 230, the NIC 250 may include a local TCB cache 280 that stores TCB information" (Paragraph [0019] – emphasis added).

However, the mere disclosure of a NIC with a provided TCP/IP offload engine and an included local TCB cache, as in Illikkal, simply fails to suggest applicant's claimed technique "wherein the CB cache is a multi-port device providing direct access to the CB cache via each port" (emphasis added), as claimed by applicant. Clearly, a NIC with a TCP/IP offload engine and a local TCB cache, as in Illikkal, simply fails to even suggest that "the CB cache is a multi-port device" (emphasis added), as claimed by applicant.

With respect to independent Claim 21, the Examiner has relied on Fig. 2 from the Illikkal reference to make a prior art showing of applicant's claimed "multi-port SRAM control block (CB) cache in a transport offload engine (TOE)... wherein the CB cache provides TOE clients direct access to the CB cache via each port."

Again, as argued hereinabove, applicant respectfully asserts that the mere disclosure of a NIC with a provided TCP/IP offload engine and an included local TCB cache, as in Illikkal, simply fails to suggest applicant's claimed "multi-port SRAM control block (CB) cache in a transport offload engine (TOE)... wherein the CB cache provides TOE clients direct access to the CB cache via each port" (emphasis added), as claimed by applicant. Clearly, a NIC with a TCP/IP offload engine and a local TCB cache, as in Illikkal, simply fails to even suggest "a multi-port SRAM control block (CB) cache" (emphasis added), as claimed by applicant.

Furthermore, with respect to independent Claims 1 and 11, the Examiner has relied on Paragraphs [0038]-[0044] from the Illikkal reference to make a prior art showing of applicant's claimed technique "wherein a first port is dedicated to transmit logic within the TOE."

Applicant respectfully asserts that the excerpt relied upon by the Examiner merely teaches "a method of send packet processing according to some embodiments" (Paragraph [0038]) and that "[b]ased on the send packet, it is predicted at 904 that a connection will subsequently have a receive packet to be processed" (Paragraph [0039]). Furthermore, the excerpt teaches that "the offload engine 270 may pre-fetch the TCB information and store it in the TCB cache 280 as scheduled" (Paragraph [0043]).

However, merely teaching a method of send packet processing where, based on the send packet, it is predicted that a connection will subsequently have a receive packet to be processed, in addition to the offload engine pre-fetching the TCB information and storing it in the TCB cache, as in Illikkal, simply fails suggest applicant's claimed technique "wherein a first port is dedicated to transmit logic within the TOE" (emphasis added), as claimed by applicant. Clearly, the offload engine pre-fetching the TCB information and storing it in the TCB cache, as in Illikkal, simply fails suggest that "a first port is dedicated to transmit logic within the TOE" (emphasis added), where "the CB cache is a multi-port device" (emphasis added), in the context as claimed by applicant.

Furthermore, with respect to independent Claims 1 and 11, the Examiner has relied on Paragraphs [0030]-[0036] and Fig. 7 from the Illikkal reference to make a prior art showing of applicant's claimed technique "wherein a second port is dedicated to receive logic within the TOE."

Applicant respectfully asserts that the excerpt and figure from the Illikkal reference relied upon by the Examiner merely discloses that "the server 200 may receive the packet at the offload engine 270 or the host processor 210" (Paragraph [0030] – emphasis added) and that "[b]ased on the receive packet, it is predicted at 704 that a connection will subsequently have a send packet to be processed" (Paragraph [0031]). Furthermore, the excerpt teaches that "the offload engine 270 might pre-fetch the TCB information and store it in the TCB cache 280 as scheduled" (Paragraph [0035]).

However, merely teaching that the server may receive the packet and that, based on the receive packet, predicting that a connection will subsequently have a send packet to be processed, in addition to the offload engine pre-fetching the TCB information and storing it in the TCB cache, as in Illikkal, simply fails to teach or suggest applicant's claimed technique "wherein a second port is dedicated to receive logic within the TOE" (emphasis added), as claimed by applicant. Clearly, the offload engine pre-fetching the TCB information and storing it in the TCB cache, as in Illikkal, simply fails suggest that "a second port is dedicated to receive logic within the TOE" (emphasis added), where "the CB cache is a multi-port device" (emphasis added), in the context as claimed by applicant.

Furthermore, with respect to independent Claim 21, it appears that the Examiner has relied on Paragraphs [0030]-[0036] and [0038]-[0044], in addition to Fig. 7 from the Illikkal reference to make a prior art showing of applicant's claimed "mechanism for dedicating ports to time critical TOE clients."

Applicant respectfully asserts that the excerpts and figure from the Illikkal reference relied upon by the Examiner merely disclose that "the server 200 may receive

the packet at the offload engine 270 or the host processor 210” (Paragraph [0030]), that “the server 900 may be ready to send the packet from the offload engine 270 or the host processor 210” (Paragraph [0038]), and that “the offload engine 270 might pre-fetch the TCB information and store it in the TCB cache 280 as scheduled” (Paragraph [0035]).

However, disclosing that the server may receive the packet at the offload engine, that the server may be ready to send the packet from the offload engine, and that the offload engine pre-fetches the TCB information and stores it in the TCB cache, as in Illikkal, simply fails to teach or suggest applicant’s claimed “mechanism for dedicating ports to time critical TOE clients” (emphasis added), as claimed by applicant.

The Examiner is reminded that a claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described in a single prior art reference. *Verdegaal Bros. v. Union Oil Co. Of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Moreover, the identical invention must be shown in as complete detail as contained in the claim. *Richardson v. Suzuki Motor Co.* 868 F.2d 1226, 1236, 9USPQ2d 1913, 1920 (Fed. Cir. 1989). The elements must be arranged as required by the claim.

This criterion has simply not been met by the above reference excerpts, as noted above.

Applicant further notes that the prior art is also deficient with respect to the dependent claims. For example, with respect to Claim 4 et al., the Examiner has relied on Fig. 2 from the Illikkal reference to make a prior art showing of applicant’s claimed “dedicating a port for host transfers of the CB entries” (as amended).

Applicant respectfully asserts that Fig. 2 from the Illikkal reference, relied upon by the Examiner, merely discloses that “[t]he server 200 includes a host processor 210” (Paragraph [0014] – emphasis added), and that “the host processor 210 may perform TCP processing (e.g., when a packet is received or is to be sent)” (Paragraph [0016] –

emphasis added). However, disclosing the server includes a host processor that may perform TCP processing, as in Illikkal, simply fails to suggest “dedicating a port for host transfers of the CB entries” (emphasis added), as claimed by applicant. Clearly, the mere disclosure that the server includes a host processor that may perform TCP processing, as in Illikkal, simply fails to even suggest “dedicating a port for host transfers of the CB entries” (emphasis added), especially where “the CB cache is a multi-port device providing direct access to the CB cache via each port,” in the context as claimed by applicant

With respect to Claim 8 et al., the Examiner has rejected the same under 35 U.S.C. 103(a) as being unpatentable over Illikkal, in view of Banerjee et al. (U.S. Patent Publication No. 2002/0165992). Specifically, the Examiner has relied on Paragraphs [0020]-[0021] (excerpted below) from the Banerjee reference to make a prior art showing of applicant’s claimed “comparing the received network packet hash value with the hash values in the hash reference table” (as amended).

“For example, WEB requests may be assigned a higher priority than other connections. Therefore, the listening socket for port 80, the WEB server port number, is assigned a high priority in order to deliver high performance. Thus, the PCB associated with this connection will be stored in the PCB cache.

According to the present invention, when a packet is received, its associated connection and socket are determined. The PCB cache is first searched to try to locate the PCB associated with this socket. If the PCB is not found in the PCB cache, the hash table or linked list is then searched. If the packet is transmitted through a connection having an associated socket which is designated as having a high priority, the PCB will be stored in the PCB cache, thus, providing a faster response. For all lower priority sockets, the associated PCBs will be stored and found in the hash table.”; (Paragraphs [0020]-[0021] - emphasis added)

Applicant respectfully asserts that the excerpt relied upon by the Examiner merely discloses that “[i]f the PCB is not found in the PCB cache, the hash table or linked list is then searched” and that “[f]or all lower priority sockets, the associated PCBs will be stored and found in the hash table” (emphasis added).

However, merely teaching that if the protocol control block (PCB) is not found in the PCB cache, the hash table is then searched, in addition to teaching that for all lower priority sockets, the associated PCBs will be stored and found in the hash table, as in Banerjee, simply fails to suggest “comparing the received network packet hash value with the hash values in the hash reference table” (emphasis added), as claimed by applicant. Clearly, searching the hash table, in addition to storing and finding associated PCBs in a hash table, as in Banerjee, simply fails to even suggest “comparing the received network packet hash value with the hash values in the hash reference table” (emphasis added), as claimed by applicant.

Further, with respect to Claim 22, the Examiner has also rejected the same under 35 U.S.C. 103(a) as being unpatentable over Illikkal, in view of Banerjee. Specifically, the Examiner has relied on the Abstract and Paragraphs [0019]-[0021] (excerpted below) from the Banerjee reference to make a prior art showing of applicant's claimed “mechanism for arbitrating access between TOE clients sharing a common port in a multi-port device based on a priority” (as amended).

“A method, system, and product are described for improving the performance of a TCP connection. A cache is established for storing protocol control blocks (PCBs). The protocol control blocks are associated with sockets which were created in response to TCP connections. PCBs are stored in the cache. PCBs may be prioritized. High priority PCBs are stored in the PCB cache while the low priority PCBs are stored in a linked list in a hash table.” (Abstract)

“A priority may be assigned to each socket, and thus to the PCB associated with the socket. For example, each socket may include a special socket option, SO_PCBCACHE, wherein the socket's priority may be stored. The PCBs associated with high priority sockets are stored in the PCB cache, while the PCBs associated with low priority sockets are stored in a linked list or hash table. In this manner, for high priority sockets, the retrieval time is reduced for the PCBs associated with these high priority sockets because the PCB cache is searched first and the PCB cache is the fastest level of memory.

For example, WEB requests may be assigned a higher priority than other connections. Therefore, the listening socket for port 80, the WEB server port number, is assigned a high priority in order to deliver high performance. Thus, the PCB associated with this connection will be stored in the PCB cache.

According to the present invention, when a packet is received, its associated connection and socket are determined. The PCB cache is first searched to try to locate the PCB associated with this socket. If the PCB is not found in the PCB cache, the hash table or linked list is then searched. If the packet is transmitted through a connection having an associated socket which is designated as having a high priority, the PCB will be stored in the PCB cache, thus, providing a faster response. For all lower priority sockets, the associated PCBs will be stored and found in the hash table.” (Paragraphs [0019]-[0021] - emphasis added)

Applicant respectfully asserts that the excerpts relied upon by the Examiner merely teach that “[a] priority may be assigned to each socket, and thus to the PCB associated with the socket” and that “[t]he PCBs associated with high priority sockets are stored in the PCB cache” (emphasis added). Further, the excerpts teach that “WEB requests may be assigned a higher priority than other connections” and “[i]f the packet is transmitted through a connection having an associated socket which is designated as having a high priority, the PCB will be stored in the PCB cache, thus, providing a faster response” (emphasis added).

However, merely disclosing that a priority may be assigned to each socket, where PCBs associated with high priority sockets are stored in the PCB cache, as in Banerjee, simply fails to teach or suggest “arbitrating access between TOE clients sharing a common port in a multi-port device based on a priority” (emphasis added), as claimed by applicant. Clearly, storing PCBs associated with high priority sockets in the PCB cache, as in Banerjee, simply fails to teach a “arbitrating access between TOE clients sharing a common port in a multi-port device based on a priority” (emphasis added), as claimed by applicant.

Further, with respect to Claim 22, the Examiner has again relied on the Abstract and Paragraphs [0019]-[0021] (excerpted above) from the Banerjee reference to make a prior art showing of applicant’s claimed technique “wherein time critical TOE clients are assigned a higher priority.” Specifically, the Examiner has argued that “certain ‘requests’, such as WEB requests, may be assigned a higher priority than other connections.”

Applicant respectfully disagrees and asserts that the excerpt from Banerjee relied upon by the Examiner merely teaches that “[f]or example, WEB requests may be assigned a higher priority than other connections” and “[t]herefore, the listening socket for port 80, the WEB server port number, is assigned a high priority in order to deliver high performance” (Paragraph [0021] – emphasis added). Further, the excerpts teach that “[i]f the packet is transmitted through a connection having an associated socket which is designated as having a high priority, the PCB will be stored in the PCB cache, thus, providing a faster response” (Paragraph [0021] – emphasis added).

However, merely assigning a higher priority to the listening socket for the WEB server port number in order to deliver high performance, in addition to storing PCB in the PCB cache if the packet is transmitted through a connection having an associated socket which is designated as having a high priority to provide a faster response, as in Banerjee, simply fails even suggest “time critical TOE clients,” much less applicant’s claimed technique “wherein time critical TOE clients are assigned a higher priority” (emphasis added), as claimed by applicant.

With respect to the rejection under 35 U.S.C. 102(e), the foregoing anticipation criterion has simply not been met by the above reference excerpt(s), as noted above.

With respect to the rejection under 35 U.S.C. 103(a), to establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art and not based on applicant’s disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed.Cir.1991).

Applicant respectfully asserts that at least the third element of the *prima facie* case of obviousness has not been met, since the prior art excerpts, as relied upon by the Examiner, fail to teach or suggest all of the claim limitations, as noted above.

Thus, a notice of allowance or a proper prior art showing of all of applicant's claim limitations, in combination with the remaining claim elements, is respectfully requested.

To this end, all of the independent claims are deemed allowable. Moreover, the remaining dependent claims are further deemed allowable, in view of their dependence on such independent claims.

In the event a telephone conversation would expedite the prosecution of this application, the Examiner may reach the undersigned at (408) 505-5100. The Commissioner is authorized to charge any additional fees or credit any overpayment to Deposit Account No. 50-1351 (Order No. NVIDP342).

Respectfully submitted,
Zilka-Kotab, PC

/KEVINZILKA/

Kevin J. Zilka
Registration No. 41,429

P.O. Box 721120
San Jose, CA 95172-1120
408-505-5100